

Robustification of Elliott's on-line EM algorithm for HMMs

Christina Erlwein and Peter Ruckdeschel

Fraunhofer ITWM
 Department of Financial Mathematics
 Fraunhofer-Platz 1, D-67663 Kaiserslautern
 Christina.Erlwein@itwm.fraunhofer.de
 Peter.Ruckdeschel@itwm.fraunhofer.de

April 9, 2013

Abstract

In this paper, we establish a robustification of an on-line algorithm for modelling asset prices within a hidden Markov model (HMM). In this HMM framework, parameters of the model are guided by a Markov chain in discrete time, parameters of the asset returns are therefore able to switch between different regimes. The parameters are estimated through an on-line algorithm, which utilizes incoming information from the market and leads to adaptive optimal estimates. We robustify this algorithm step by step against additive outliers appearing in the observed asset prices with the rationale to better handle possible peaks or missings in asset returns.

Keywords: Robustness, HMM, Additive outlier, Asset pricing

1 Introduction

Realistic modelling of financial time series from various markets (stocks, commodities, interest rates etc.) in recent years often is achieved through hidden Markov or regime-switching models. One major advantage of regime-switching models is their flexibility to capture switching market conditions or switching behavioural aspects of market participants resulting in a switch in the volatility or mean value.

Regime-switching models were first applied to issues in financial markets through Hamilton [1989], where he established a Markov switching AR-model to model the GNP of the U.S. His results show promising effects of including possible regime-switches into the characterisation of a financial time series. A lot of further approaches to use regime-switching models for financial time series followed, e.g. switching ARCH or switching GARCH models (see for example Cai [1994] and Gray [1996]), amongst many other applications.

Various algorithms and methods for statistical inference are applied within these model set-ups, including as famous ones as the Baum-Welch algorithm and Viterbi's algorithm for an estimation of the optimal state sequence. HMMs in Finance, both in continuous and in discrete time often utilise a filtering technique which was developed by Elliott [1994]. Adaptive filters are derived for processes of the Markov chain (jump process, occupation time process and auxiliary processes) which are in turn used for recursive optimal parameter estimates of the model parameters. This filter-based Expectation-Maximization (EM) algorithm leads to an on-line estimation of model parameters. Our model set-up is based on Elliott's filtering framework.

This HMM can be applied to questions, which arise in asset allocation problems. An investor typically has to decide, how much of his wealth shall be invested into which asset or asset class and when to optimally restructure a portfolio. Asset allocation problems were examined in a regime-switching setting by Ang and Bekaert [2002], where high volatility and high correlation regimes of asset returns were discovered. Guidolin and Timmermann [2007] presented an asset allocation problem within a regime-switching model and found four different possible underlying market regimes. A paper by Sass and Haussmann [2004] derives optimal trading strategies and filtering techniques in a continuous-time regime-switching model set up. Optimal portfolio choices were also discussed in Elliott and van der Hoek [1997] and Elliott and Hinz [2003] amongst others. Here, Markowitz's

famous mean-variance approach (see Markowitz [1952]) is transferred into an HMM and optimal weights are derived. A similar Markowitz based approach within an HMM was developed in Erlwein et al. [2011], where optimal trading strategies for portfolio decisions with two asset classes are derived. Trading strategies are developed herein to find optimal portfolio decision for an investment in either growth or value stocks. Elliott's filtering technique is utilised to predict asset returns.

However, most of the optimal parameter estimation techniques for HMMs in the literature only lead to reasonable results, when the market data set does not contain significant outliers. The handling of outliers is an important issue in many financial models, since market data might be unreliable at times or high peaks in asset returns, which might occur in the market from time to time shall be considered separately and shall not negatively influence the parameter estimation method. In general, higher returns in financial time series might belong to a separate regime within an HMM. This flexibility is already included in the model set-up. However, single outliers, which are not typical for any of the regimes considered, shall be handled with care, a separate regime would not reflect the abnormal data point. In this paper, we will develop a robustification of Elliot's filter-based EM-algorithm. In section 2 we will set the HMM framework, which is applied (either in a one- or multi-dimensional setting) to model asset or index returns. The general filtering technique is described in section 3. The asset allocation problem which clarifies the effect outliers can have on the stability of the filters is developed in section 4. Section 5 then states the derivation of a robustification for various steps in the filter equations. The robustification of a reference probability measure is derived as well as a robust version of the filter-based EM-algorithm. An application of the robust filters is shown in section 6 and section 7 finishes our work with some conclusions and possible future applications.

2 Hidden Markov model framework for asset returns

For our problem setting we first review a filtering approach for a hidden Markov model in discrete time which was developed by Elliott [1994]. The logarithmic returns of a stock or an index follow the dynamics of the observation process y_k , which can be interpreted as a discretized version of the Geometric Brownian motion, which is a standard process to model stock returns. The underlying hidden Markov chain \mathbf{x}_k cannot be directly observed. The parameters of the observation process are governed by the Markov chain and are therefore able to switch between regimes over time.

We work under a probability space (Ω, \mathcal{F}, P) under which \mathbf{x}_k is a homogeneous Markov chain with finite state space $I = \{1, \dots, N\}$ in discrete time ($k = 0, 1, 2, \dots$). Let the state space of \mathbf{x}_k be associated with the canonical basis $\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_N\} \in \mathbb{R}^N$ with $\mathbf{e}_i = (0, \dots, 0, 1, 0, \dots, 0)^\top \in \mathbb{R}^N$. The initial distribution of \mathbf{x}_0 is known and $\Pi = (\pi_{ji})$ is the transition probability matrix with $\pi_{ji} = P(\mathbf{x}_{k+1} = \mathbf{e}_j | \mathbf{x}_k = \mathbf{e}_i)$. Let $\mathcal{F}_k^{\mathbf{x}_0} = \sigma\{\mathbf{x}_0, \dots, \mathbf{x}_k\}$ be the σ -field generated by $\mathbf{x}_0, \dots, \mathbf{x}_k$ and let $\mathcal{F}_k^{\mathbf{x}}$ be the complete filtration generated by $\mathcal{F}_k^{\mathbf{x}_0}$. Under the real world probability measure P , the Markov chain \mathbf{x} has the dynamics

$$\mathbf{x}_{k+1} = \Pi \mathbf{x}_k + \mathbf{v}_{k+1} \quad (2.1)$$

where $\mathbf{v}_{k+1} := \mathbf{x}_{k+1} - \Pi \mathbf{x}_k$ is a martingale increment (see Theorem in Elliott [1994]).

The Markov chain \mathbf{x}_k is "hidden" in the log returns y_{k+1} of the stock price S_k . Our observation process is given by

$$y_{k+1} = \ln \frac{S_{k+1}}{S_k} = f(\mathbf{x}_k) + \sigma(\mathbf{x}_k) w_{k+1} \quad (2.2)$$

where \mathbf{x}_k has finite state space and w_k 's constitute a sequence of i.i.d. random variables independent of \mathbf{x} . The real-valued process y can be re-written as

$$y_{k+1} = \langle \mathbf{f}, \mathbf{x}_k \rangle + \langle \boldsymbol{\sigma}, \mathbf{x}_k \rangle w_{k+1}. \quad (2.3)$$

Note that $\mathbf{f} = (f_1, f_2, \dots, f_N)^\top$ and $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \dots, \sigma_N)^\top$ are vectors, furthermore $f(\mathbf{x}_k) = \langle \mathbf{f}, \mathbf{x}_k \rangle$ and $\sigma(\mathbf{x}_k) = \langle \boldsymbol{\sigma}, \mathbf{x}_k \rangle$, where $\langle \mathbf{b}, \mathbf{c} \rangle$ denotes the Euclidean scalar product in \mathbb{R}^N of the vectors \mathbf{b} and \mathbf{c} . We assume $\sigma_i \neq 0$. Let \mathcal{F}_k^y be the filtration generated by the $\sigma(y_1, y_2, \dots, y_k)$ and $\mathcal{F}_k = \mathcal{F}_k^{\mathbf{x}} \vee \mathcal{F}_k^y$ is the global

filtration.

The following theorem (Elliott [1994]) states that the dynamics of the underlying Markov chain can be described by martingale differences.

3 Essential Steps in Elliott's Algorithm

3.1 Change of Measure

A widely used concept in filtering applications, going back to Zakai [1969] for stochastic filtering, is a change of probability measure technique. A measure change to a reference measure \bar{P} is applied here, under which filters for the Markov chain and related processes are derived. Under \bar{P} , the underlying Markov chain still has the dynamics $\mathbf{x}_{k+1} = \mathbf{P}\mathbf{x}_k + \mathbf{v}_{k+1}$ but is independent of the observation process and the observations y_k are $\mathcal{N}(0, 1)$ i.i.d. random variables.

Following the change of measure technique which was outlined in Elliott et al. [1995] the adaptive filters for the Markov chain and related processes are derived under this “idealised” measure \bar{P} . Changing back to the real world is done by constructing P from \bar{P} through the Radon-Nikodým derivative $\left. \frac{dP}{d\bar{P}} \right|_{\mathcal{F}_k} = \Lambda_k$. To construct Λ_k we define the process λ_l

$$\lambda_l := \frac{\phi\left[\sigma(\mathbf{x}_{l-1})^{-1}(y_l - f(\mathbf{x}_{l-1}))\right]}{\sigma(\mathbf{x}_{l-1})\phi(y_l)} \quad (3.1)$$

where $\phi(z)$ is the probability density function of a standard normal random variable Z and set $\Lambda_k := \prod_{l=1}^k \lambda_l$, $k \geq 1$, $\Lambda_0 = 1$. Under P the sequence of variables w_1, w_2, \dots , is a sequence of i.i.d. standard normals, where we have $w_{k+1} = \sigma(\mathbf{x}_k)^{-1}(y_{k+1} - f(\mathbf{x}_k))$.

3.2 Filtering for general adapted processes

The general filtering techniques and the filter equations which were established by Elliott [1994] for Markov chains observed in Gaussian noise are stated in this subsection. This filter-based EM-algorithm is adaptive, which enables fast calculations and filter updates. Our robustification partly keeps this adaptive structure of the algorithm, although the recursivity cannot be kept completely.

In general, filters for four types of processes related to the Markov chain, namely the state space process, the jump process, the occupation time process and auxiliary processes including terms of the observation process are derived. Information on these processes can be filtered out from our observation process and can in turn be used to find optimal parameter estimates.

To determine the expectation of any \mathcal{F} -adapted stochastic process H given the filtration \mathcal{F}_k^y , consider the reference probability measure \bar{P} defined as $P(A) = \int_A \Lambda d\bar{P}$. From Bayes' theorem a filter for any adapted process H is given by $E[H_k | \mathcal{F}_k^y] = \bar{E}[H_k \Lambda_k | \mathcal{F}_k^y] / \bar{E}[\Lambda_k | \mathcal{F}_k^y]$. We define $\eta_k(H_k) := \bar{E}[H_k \Lambda_k | \mathcal{F}_k^y]$, so that $E[H_k | \mathcal{F}_k^y] = \eta_k(H_k) / \eta_k(1)$. A recursive relationship between $\eta_k(H_k)$ and $\eta_{k-1}(H_{k-1})$ has to be found, where $\eta_0(H_0) = E[H_0]$. However, a recursive formula for the term $\eta_{k-1}(H_{k-1} \mathbf{x}_{k-1})$ is found. To relate $\eta_k(H_k)$ and $\eta_k(H_k \mathbf{x}_k)$ we note that with $\langle \mathbf{1}, \mathbf{x}_k \rangle = 1$

$$\langle \mathbf{1}, \eta_k(H_k \mathbf{x}_k) \rangle = \eta_k(H_k \langle \mathbf{1}, \mathbf{x}_k \rangle) = \eta_k(H_k). \quad (3.2)$$

Therefore

$$E[H_k | \mathcal{F}_k^y] = \frac{\langle \mathbf{1}, \eta_k(H_k \mathbf{x}_k) \rangle}{\langle \mathbf{1}, \eta_k(\mathbf{x}_k) \rangle}. \quad (3.3)$$

A general recursive filter for adapted processes was derived by Elliott [1994]. Suppose H_l is a scalar $\mathcal{F} = \sigma((\mathbf{x}_t, Y_t)_t)$ -adapted process, H_0 is \mathcal{F}_0^x measurable and $H_l = H_{l-1} + a_l + \langle b_l, \mathbf{v}_l \rangle + g_l f(y_l)$, where a, b

and g are \mathcal{F} -predictable, f is a scalar-valued function and $\mathbf{v}_l = \mathbf{x}_l - \Pi \mathbf{x}_{l-1}$. A recursive relation for $\eta_k(H_k \mathbf{x}_k)$ is given by

$$\begin{aligned} \eta_k(H_k \mathbf{x}_k) = & \sum_{i=1}^N \Gamma^i(y_k) [\langle \mathbf{e}_i, \eta_{k-1}(H_{k-1} \mathbf{x}_{k-1}) \rangle \Pi \mathbf{e}_i \\ & + \langle \mathbf{e}_i, \eta_{k-1}(a_k \mathbf{x}_{k-1}) \rangle \Pi \mathbf{e}_i \\ & + (\text{diag}(\Pi \mathbf{e}_i) - (\Pi \mathbf{e}_i)(\Pi \mathbf{e}_i)') \eta_{k-1}(b_k \langle \mathbf{e}_i, \mathbf{x}_{k-1} \rangle) \\ & + \eta_{k-1}(g_k \langle \mathbf{e}_i, \mathbf{x}_{k-1} \rangle) f(y_k) \Pi \mathbf{e}_i] \end{aligned} \quad (3.4)$$

Here, for any column vectors \mathbf{z} and \mathbf{y} , $\mathbf{z}\mathbf{y}'$ denotes the rank-one (if $\mathbf{z} \neq \mathbf{0}$ and $\mathbf{y} \neq \mathbf{0}$) matrix $\mathbf{z}\mathbf{y}'$. The term $\Gamma^i(y_k)$ denotes the component-wise Radon-Nikod m derivative λ_k^i :

$$\Gamma^i(y_k) = \phi\left(\frac{y_k - f_i}{\sigma_i}\right) / \sigma_i \phi(y_k)$$

Now, filters for the state of the Markov chain as well as for three related processes: the jump process, the occupation time process and auxiliary processes of the Markov chain are derived. These processes can be characterised as special cases of the general process H_l .

The estimator for the state \mathbf{x}_k is derived from $\eta_k(H_k \mathbf{x}_k)$ by setting $H_k = H_0 = 1$, $a_k = 0$, $b_k = 0$ and $g_k = 0$. This implies that

$$\eta_k(\mathbf{x}_k) = \sum_{i=1}^N \Gamma^i(y_k) \langle \mathbf{e}_i, \eta_{k-1}(\mathbf{x}_{k-1}) \rangle \Pi \mathbf{e}_i . \quad (3.5)$$

The first related process is the number of jumps of the Markov chain \mathbf{x}_k from state \mathbf{e}_r to state \mathbf{e}_s in time k , $J_k^{(sr)} = \sum_{l=1}^k \langle \mathbf{x}_{l-1}, \mathbf{e}_r \rangle \langle \mathbf{x}_l, \mathbf{e}_s \rangle$. Setting $H_k = J_k^{(sr)}$, $H_0 = 0$, $a_k = \langle \mathbf{x}_{k-1}, \mathbf{e}_r \rangle \pi_{sr}$, $b_k = \langle \mathbf{x}_{k-1}, \mathbf{e}_r \rangle \mathbf{e}_s'$ and $g_k = 0$ in equation (3.4) we get

$$\begin{aligned} \eta_k(J_k^{sr} \mathbf{x}_k) = & \sum_{i=1}^N \Gamma^i(y_k) \langle \eta_{k-1}(J_{k-1}^{sr} \mathbf{x}_{k-1}), \mathbf{e}_i \rangle \Pi \mathbf{e}_i \\ & + \Gamma^r(y_k) \eta_{k-1}(\langle \mathbf{x}_{k-1}, \mathbf{e}_r \rangle) \pi_{sr} \mathbf{e}_s . \end{aligned} \quad (3.6)$$

The second process $O_k^{(r)}$ denotes the occupation time of the Markov process \mathbf{x} , which is the length of time \mathbf{x} spent in state r up to time k . Here, $O_k^r = \sum_{l=1}^k \langle \mathbf{x}_{l-1}, \mathbf{e}_r \rangle = O_{k-1}^r + \langle \mathbf{x}_{k-1}, \mathbf{e}_r \rangle$. We set $H_k = O_k^r$, $H_0 = 0$, $a_k = \langle \mathbf{x}_{k-1}, \mathbf{e}_r \rangle$, $b_k = 0$ and $g_k = 0$ in equation (3.4) to obtain

$$\begin{aligned} \eta_k(O_k^r \mathbf{x}_k) = & \sum_{i=1}^N \Gamma^i(y_k) \langle \eta_{k-1}(O_{k-1}^r \mathbf{x}_{k-1}), \mathbf{e}_i \rangle \Pi \mathbf{e}_i \\ & + \Gamma^r(y_k) \langle \eta_{k-1}(\mathbf{x}_{k-1}), \mathbf{e}_r \rangle \Pi \mathbf{e}_r . \end{aligned} \quad (3.7)$$

Finally, consider the auxiliary process $T_k^r(g)$, which occur in the maximum likelihood estimation of model parameters. Specifically, $T_k^{(r)}(g) = \sum_{l=1}^k \langle \mathbf{x}_{l-1}, \mathbf{e}_r \rangle g(y_l)$, where g is a function of the form $g(y) = y$ or $g(y) = y^2$. We apply formula (3.4) and get

$$\begin{aligned} \eta_k(T_k^r(g) \mathbf{x}_k) = & \sum_{i=1}^N \Gamma^i(y_k) \langle \eta_{k-1}(T_{k-1}^r(g(y_{k-1})) \mathbf{x}_{k-1}), \mathbf{e}_i \rangle \Pi \mathbf{e}_i \\ & + \Gamma^r(y_k) \langle \eta_{k-1}(\mathbf{x}_{k-1}), \mathbf{e}_r \rangle g(y_k) \Pi \mathbf{e}_r . \end{aligned} \quad (3.8)$$

The recursive optimal estimates of J , O and T can be calculated using equation (3.2).

3.3 Filter-based EM-algorithm

The derived adapted filters for processes of the Markov chain can now be utilised to derive optimal parameter estimates through a filter-based EM-algorithm. The set of parameters ρ , which determines the regime-switching model is

$$\rho = \{\pi_{ji}, 1 \leq i, j \leq N, f_i, \sigma_i, 1 \leq i \leq N\}. \quad (3.9)$$

Initial values for the EM algorithm are assumed to be given. Starting from these values updated parameter estimates are derived which maximise the conditional expectation of the log-likelihoods. The M-step of the algorithm deals with maximizing the following likelihoods:

M-Step

- The likelihood in the global \mathcal{F} -model is given by

$$\log \Lambda_t(\sigma, f; (\mathbf{x}_s, y_s)_{s \leq t}) = -\frac{1}{2} \sum_{s=1}^t \left(\log \langle \sigma, \mathbf{x}_{s-1} \rangle + \frac{(y_s - \langle f, \mathbf{x}_{s-1} \rangle)^2}{\langle \sigma, \mathbf{x}_s \rangle} \right)$$

- In the \mathcal{F}_y -model, where the Markov chain is not observed, we obtain

$$\begin{aligned} L_t(\sigma, f; (y_s)_{s \leq t}) &= E[\log \Lambda_t(\sigma, f; (\mathbf{x}_s, y_s)_{s \leq t}) \mid \mathcal{F}_t] = \\ &= -\frac{1}{2} \sum_{k=1}^N \left(\log \sigma_k \hat{O}_t^k + (\hat{T}_t^k(y^2) - 2\hat{T}_t^k(y)f_k + \hat{O}_t^k f^2) / \sigma_k^2 \right) \end{aligned} \quad (3.10)$$

The maximum likelihood estimates of the model parameters can be expressed through the adapted filters. Whenever new information is available on the market, the filters are updated and, respectively, updated parameter estimates can be obtained.

Theorem 3.1 (Optimal parameter estimates) Write $\hat{H}_k = E[H_k | \mathcal{F}_k^y]$ for any adapted process H . With \hat{J} , \hat{O} and \hat{T} denoting the best estimates for the processes J , O and T , respectively, the optimal parameter estimates $\hat{\pi}_{ji}$, \hat{f}_i and $\hat{\sigma}_i$ are given by

$$\hat{\pi}_{ji} = \frac{\hat{J}_k^{ji}}{\hat{O}_k^i} = \frac{\eta_k(J_k^{ji})}{\eta_k(O_k^i)} \quad (3.11)$$

$$\hat{f}_i = \frac{\hat{T}_k^{(i)}}{\hat{O}_k^{(i)}} = \frac{\eta(T^{(i)}(y))_k}{\eta(O^{(i)})_k} \quad (3.12)$$

$$\hat{\sigma}_i = \sqrt{\frac{\hat{T}_k^{(i)}(y^2) - 2\hat{f}_i \hat{T}_k^{(i)}(y) + \hat{f}_i^2 \hat{O}_k^{(i)}}{\hat{O}_k^{(i)}}}. \quad (3.13)$$

PROOF The derivation of the optimal parameter estimates can be found in Elliott et al. [1995].

////

Summary The filter-based EM-algorithm runs in batches of n data points (n typically equals to a minimum of ten up to a maximum of fifty) over the given time series. The parameters are updated at the end of each batch. Elliott's Algorithm comprises the following steps

- (0) Find suitable starting values for Π and f , σ .
- (RN) Determine the RN-derivative for the measure change to \bar{P} .
- (E) Recursively, compute filters \hat{J}_t^{ji} , \hat{O}_t^i , and $\hat{T}_t^i(g)$.
- (M1) Obtain ML-estimators $\hat{f} = (\hat{f}_1, \dots, \hat{f}_N)$ and $\hat{\sigma} = (\hat{\sigma}_1, \dots, \hat{\sigma}_N)$.
- (M2) Obtain ML-estimators Π .
- (Rec) Go to (RN) to compute the next batch.

4 Outliers in Asset Allocation Problem

4.1 Outliers in General

In the following sections we derive a robustification of the Algorithm 3.3 to stabilize it in the presence of outliers in the observation process. To this end let us discuss what makes an observation an outlier. First of all, outliers are exceptional events, occurring **rarely**, say with probability 5%–10%. Rather than captured by usual randomness, i.e., by some distributional model, they belong to what Knight [1921] refers to *uncertainty*: They are **uncontrollable**, of **unknown** distribution, **unpredictable**, their distribution may change from observation to observation, so they are non-recurrent and do not form an additional state, so cannot be used to enhance predictive power, and, what makes their treatment difficult, they often cannot be told with certainty from ideal observations.

Still, the majority of the observations in a realistic sample should resemble an ideal (distributional) setting closely, otherwise the modeling would be questionable. Here we understand closeness as in a distributional sense, as captured, e.g., by goodness-of-fit distances like Kolmogorov, total variation or Hellinger distance. More precisely, ideally, this closeness should be compatible to the usual convergence mode of the Central Limit Theorem, i.e., with weak topology. In particular, closeness in moments is incompatible with this idea.

Topologically speaking, one would most naturally use balls around a certain element, i.e., the set of all distributions with a suitable distance no larger than some given radius $\varepsilon > 0$ to the distribution assumed in the ideal model.

Conceptually, the most tractable neighborhoods are given by the so-called *Gross Error Model*, defining a neighborhood \mathcal{U} about a distribution F as the set of all distributions given by

$$\mathcal{U}_c(F, \varepsilon) = \{G \mid \exists H: G = (1 - \varepsilon)F + \varepsilon H\} \quad (4.1)$$

They can also be thought of as the set of all distributions of (*realistic*) random variables X^{re} constructed as

$$X^{\text{re}} = (1 - U)X^{\text{id}} + UX^{\text{di}} \quad (4.2)$$

where X^{id} is a random variable distributed according to the *ideal* distribution and U is an independent $\text{Bin}(1, \varepsilon)$ switching variable, which in most cases lets you see X^{id} but in some cases replaces it by some contaminating or *distorting* variable X^{di} which has nothing to do with the original situation.

4.2 Time-dependent Context: Exogenous and Endogenous Outliers

In our time dependent setup in addition to the i.i.d. situation, we have to distinguish whether the impact of an outlier is propagated to subsequent observations or not. Historically there is a common terminology due to Fox [1972], who distinguishes *innovation outliers* (or IO's) and *additive outliers* (or AO's). Non-propagating AO's are added at random to single observations, while IO's denote gross errors affecting the innovations. For consistency with literature, we use the same terms, but use them in a wider sense: *IO's* stand for general endogenous outliers entering the state layer (or the Markov chain in the present context), hence with propagated distortion. As in our Markov chain, the state space is finite, IO's are much less threatening as they are in general.

Correspondingly, wide-sense *AO's* denote general exogenous outliers which do not propagate, hence also comprise substitutive outliers or *SO's* as defined in a simple generalization of (4.2) to the state space context in equations (4.3)–(4.6).

$$Y^{\text{re}} = (1 - U)Y^{\text{id}} + UY^{\text{di}}, \quad U \sim \text{Bin}(1, r) \quad (4.3)$$

for U independent of $(X, Y^{\text{id}}, Y^{\text{di}})$ and some arbitrary distorting random variable Y^{di} for which we assume

$$Y^{\text{di}}, X \text{ independent} \quad (4.4)$$

and the law of which is arbitrary, unknown and uncontrollable. As a first step consider the set $\partial\mathcal{U}^{\text{SO}}(r)$ defined as

$$\partial\mathcal{U}^{\text{SO}}(r) = \left\{ \mathcal{L}(X, Y^{\text{re}}) \mid Y^{\text{re}} \text{ acc. to (4.3) and (4.4)} \right\} \quad (4.5)$$

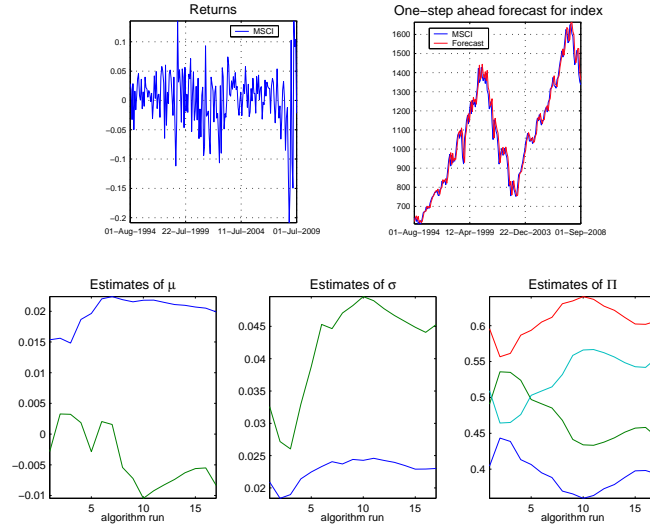


Figure 1: Optimal parameter estimates for monthly MSCI returns between 1994 and 2009

Because of condition (4.4), in the sequel we refer to the random variables Y^{re} and Y^{di} instead of their respective (marginal) distributions only, while in the common gross error model, reference to the respective distributions would suffice. Condition (4.4) also entails that in general, contrary to the gross error model, $\mathcal{L}(X, Y^{\text{id}})$ is not element of $\partial\mathcal{U}^{\text{SO}}(r)$, i.e., not representable itself as some $\mathcal{L}(X, Y^{\text{re}})$ in this neighborhood.

As corresponding (convex) neighborhood we define

$$\mathcal{U}^{\text{SO}}(r) = \bigcup_{0 \leq s \leq r} \partial\mathcal{U}^{\text{SO}}(s) \quad (4.6)$$

hence the symbol “ ∂ ” in $\partial\mathcal{U}^{\text{SO}}$, as the latter can be interpreted as the corresponding surface of this ball. Of course, $\mathcal{U}^{\text{SO}}(r)$ contains $\mathcal{L}(X, Y^{\text{id}})$. In the sequel where clear from the context we drop the superscript SO and the argument r .

Due to their different nature, as a rule, IO’s and AO’s require different policies: As AO’s are exogenous, we would like to damp their effect, while when there are IO’s, something has happened in the system, so the usual goal will be to track these changes as fast as possible.

4.3 Evidence for Robustness Issue in Asset Allocation

In this section we examine the robustness of the filter and parameter estimation technique. The filter technique is implemented and applied to monthly returns of MSCI index between 1994 and 2009. The MSCI World Index is one of the leading indices on the stock markets and a common benchmark for global stocks. The algorithm is implemented with batches of ten data points, therefore the adaptive filters are updated whenever ten new data points are available on the market. The recursive parameter estimates, which utilise this new information, are updated as well, the algorithm is self-tuning. Care has to be taken when choosing the initial values for the algorithm, since the EM-algorithm in its general form converges to a local maximum. In this implementation we choose the initial values with regard to mean and variance of the first ten data points. Figure 1 shows the original return series, optimal parameter estimates for the index returns as well as the one-step ahead forecast.

To highlight the sensitivity of the filter technique towards exogenous outliers we plant unusual high returns within the time series. Considerable SO outliers are included at time steps $t = 40, 80, 130, 140$. The optimal parameter estimation through the filter-based EM-algorithm of this data set with outliers can be seen in Figure 2. The filter still finds optimal parameter estimates, although the estimates are visibly affected by the outliers.

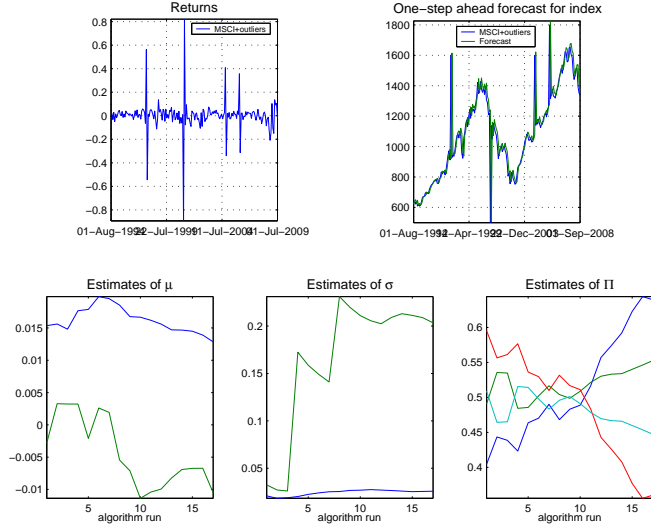


Figure 2: Optimal parameter estimates for monthly MSCI returns with planted outliers

In a third step, severe outliers are planted into the observation sequence. Data points $t = 40, 80, 130, 140$ now show severe SO outliers as can be seen from the first panel in Figure 3. The filters cannot run through any longer, optimal parameter estimates cannot be established in a setting with severe outliers.

In practice, asset or index return time series can certainly include outliers from time to time. This might be due to wrong prices in the system, but also due to very unlikely market turbulence for a short period of time. It has to be noted, that the type of outliers which we consider in this study does not characterise an additional state of the Markov chain. In the following, we develop robust filter equations, which can handle exogenous outliers.

5 Robust Statistics

To overcome effects like in Figure 3 we need more stable variants of the Elliott type filters discussed so far. This is what robust statistics is concerned with. Excellent monographs on this topic are e.g., Huber [1981], Hampel et al. [1986], Rieder [1994], Maronna et al. [2006]. This section provides necessary concepts and results from robust statistics needed to obtain the optimally-robust estimators used in this article.

5.1 Concepts of Robust Statistics

The central mathematical concepts of continuity, differentiability, or closeness to singularities may in fact serve to operationalize stability quite well already. To make these available in our context, it helps to consider a statistical procedure, i.e.; an estimator, a predictor, a filter, or a test as a function of the underlying distribution. In a parametric context, this amounts to considering functionals T mapping distributions to the parameter set Θ . An estimator will then simply be T applied to the empirical distribution \hat{F}_n . For filtering or prediction, the range of such a functional will rather be the state space but otherwise the arguments run in parallel.

For a notion of continuity, we have to specify a topology, and as in case of outliers, we use topologies essentially compatible with the weak topology. With these neighborhoods, we now may easily translate the notions of continuity, differentiability and closest singularity to this context: (Equi-)continuity is then called *qualitative robustness* [Hampel et al., 1986, Sec. 2.2 Def. 3], a differentiable functional with a bounded derivative is called *local robust*, and its derivative is called *influence function* (IF)¹, compare [Hampel et al., 1986, Sec. 2.1 Def. 1].

¹In mathematical rigor, the IF, when it exists, is the Gâteaux derivative of functional T into the direction of the tangent $\delta_x - F$. For certain properties, this notion is in fact too weak, and one has to require stronger notions like Hadamard or Fréchet differentiability; for details, see Fernholz [1983] or [Rieder, 1994, Ch. 1].

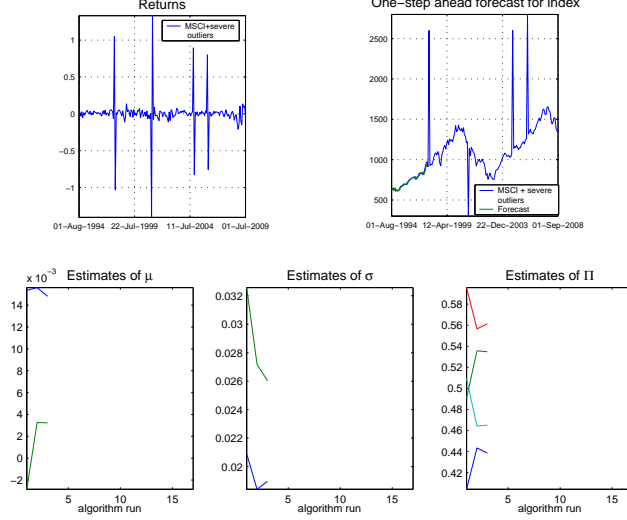


Figure 3: Filter-based EM-algorithm for observation sequence with severe outliers

The IF reflects the infinitesimal influence of a single observation on the estimator. Under additional assumptions, many of the asymptotic properties of an estimator are expressions in the IF ψ . E.g., the asymptotic variance of the estimator in the ideal model is the second moment of ψ . Infinitesimally, i.e., for $\varepsilon \rightarrow 0$, the maximal bias on \mathcal{U} is just $\sup |\psi|$, where $|\cdot|$ denotes Euclidean norm. $\sup |\psi|$ is then also called *gross error sensitivity* (GES), [Hampel et al., 1986, (2.1.13)]. Seeking robust optimality hence amounts to finding optimal IFs.

To grasp the maximal bias of a functional T on a neighborhood $\mathcal{U} = \mathcal{U}(F; \varepsilon)$ of radius ε , one considers the *max-bias curve* $\varepsilon \mapsto \sup_{G \in \mathcal{U}(F; \varepsilon)} |T(G) - T(F)|$. The singularity of this curve closest to 0 (i.e., the ideal situation of no outliers) captures its behavior under massive deviations, or its global robustness. In robust statistics, this is called *breakdown point*—the maximal radius ε the estimator can cope with without producing an arbitrary large bias, see [Hampel et al., 1986, Sec. 2.2 Def.’s 1,2] for formal definitions.

Usually, the classically optimal estimators (MLE in many circumstances) are non-robust, both locally and globally. Robust estimators on the other hand pay a certain price for this stability as expressed by an *asymptotic relative efficiency* (ARE) strictly lower than 1 in the ideal model, where ARE is the ratio of the two asymptotic (co)variances of the classically optimal estimator and its robust alternative.

To rank various robust procedures among themselves, other quality criteria are needed, though, summarizing the behavior of the procedure on a whole neighborhood, as in (4.1). A natural candidate for such a criterion is maximal MSE (maxMSE) on some neighborhood \mathcal{U} around the ideal model and, for estimation context, maximal bias (maxBias) on the respective neighborhood, or, referring to famous [Hampel, 1968, Lemma 5], trace of the ideal variance subject to a bias bound on this neighborhood. In estimation context, the respective solution are usually called OMSE (*Optimal MSE estimator*), MBRE (*Most Bias Robust Estimator*), and OBRE (*Optimally Bias Robust Estimator*)².

In our context we encounter two different situations where we want to apply robust ideas: (recursive) filtering in the (E)-step and estimation in the (M)-step. While in the former situation we only add a single new observation, which precludes asymptotic arguments, in the (M)-step, the preceding application of Girsanov’s theorem turns our situation into an i.i.d. setup, where each observation becomes (uniformly) asymptotically negligible and asymptotics apply in a standard form.

5.2 Our Robustification of the HMM: General Strategy

As a robustification of the whole estimation process in this only partially observed model would (a) lead to computationally intractable terms and (b) would drop the key feature of recursivity, we instead propose to robustify each of the steps in Elliott’s algorithm separately. Doing so, the whole procedure will be robust,

²For the terms OBRE and MBRE, see Hampel et al. [1986], while for OMSE see Ruckdeschel and Horbenko [2010].

but in general will loose robust optimality, i.e.; contrary to multiparameter maximum likelihood, the Bellmann principle does not hold for optimal robust multi-step procedures simply because optimal clipping in several steps is not the same as joint optimal clipping of all steps. Table 1 lists all essential steps in Elliot’s algorithm related to our proposed robustification approach.

Classical setting	Robust version
Initialization: Find suitable starting values for Π , f , σ and \mathbf{x}_0 .	
Build N clusters on first batches	Build $N + 1$ clusters on first batches, distribute points in outlier cluster randomly on other clusters.
Use first and second moment of each cluster as initial values for f and σ .	Use median and MAD of clusters for f and σ .
Choose Π and \mathbf{x}_0 according to cluster probabilities.	Choose Π and \mathbf{x}_0 according to cluster probabilities.
E-step: Determine RN-derivative and calculate E-Step.	
Find RN-derivative. Λ	Robustified version of Λ through suitable clipped version of λ_k .
Estimate recursive filters J_k^{ji} , O_k^i and $T_k^i(g)$.	No further robustification needed; i.e., take over J_k^{ji} , O_k^i unchanged and skip $T_k^i(g)$.
M-step 1: Obtain estimates for f and σ .	
MLE-estimates for f and σ through recursive filters O_k^i and $T_k^i(g)$.	Likelihoods re-stated; they are expressed as weighted sums of the observations y_k .
Recursive filters are substituted into likelihood.	Robustified version of MLE through asymptotic linear estimators.
Estimates updated after each batch.	Estimates updated after each batch, recursivity cannot be preserved completely.
M-step 2: Obtain ML-estimators Π .	
MLE-estimation, recursive filters J_k^{ji} and O_k^i are substituted into likelihood.	Robustification through robust version of filters. J_k^{ji} and O_k^i , no further observation y_k has to be considered.
Rec: Algorithm runs on next batch.	
Go to (RN) to compute the next batch.	Go to (RN) to compute the next batch.

Table 1: Classical algorithm setting and robustified version of each step.

5.3 Robustification of Step (0)

So far, little has been said as to the initialization even in the non-robustified setting. Basically, all we have to do is to make sure that the EM algorithm converges. In prior applications of this algorithms (Mamon et al. [2008] and Erlwein et al. [2011] amongst others), one approach was to fill Π with entries $1/N$, i.e., with uniform (and hence non-informative) distribution over all states, independent from state \mathbf{x}_0 . As to f_i and σ_i , an ad hoc approach would estimate the global mean and variance over all states and then, again in a non-informative way jitter the state-individual moments, adding independent noise to it. In our preliminary experiments, it turned

out that this naïve approach could drastically fail in the presence of outliers, so we instead propose a more sophisticated approach which can also be applied in a classical (i.e., non-robust) setting: In a first step we ignore the time dynamics and interpret our observations as realizations of a Gaussian Mixture Model, for which we use R package `mclust` (Fraley and Raftery [2002], Fraley et al. [2012]) to identify the mixture components, and for each of these, we individually determine the moments f_i and σ_i . As to Π , we again assume independence of \mathbf{x}_0 , but fill the columns according to the estimated frequencies of the mixture components. In case of the non-robust setting we would use N mixture components and for each of them determine f_i and σ_i by their ML estimators (assuming independent observations). For a robust approach, we use $N+1$ mixture components, one of them—the one with the lowest frequency—being a pure noise component capturing outliers. For each non-noise component we retain the ML estimates for f_i and σ_i . The noise component is then randomly distributed amongst the remaining components, respecting their relative frequencies prior to this redistribution.

We are aware of the fact that reassigning the putative outliers at random could be misleading in ideal situations (with no outliers) where one cluster could be split off into two but not necessarily so. Then in our strategy, the smaller offspring of this cluster would in part be reassigned to wrong other clusters, so this could still be worked on. On the other hand, this choice often works reasonably well, and as more sophisticated strategies are questions for model selection, we defer them to further work.

Based on the f_i and σ_i , for each observation j and each state i , we get weights $0 \leq w_{i,j} \leq 1$, $\sum_i w_{i,j} = 1$ for each j , representing the likelihood that observation j is in state i . For each i , again we determine robustified moment estimators f'_i , σ'_i as weighted medians and scaled weighted MADs (medians of absolute deviations).

Weighted Medians And MADs: For weights $w_j \geq 0$ and observations y_j , the weighted median $m = m(y, w)$ is defined as $m = \operatorname{argmin}_f \sum_j w_j |y_j - f|$, and with $y'_j = |y_j - m|$, the scaled weighted MAD $s = s(y, w)$ is defined as $s = c^{-1} \operatorname{argmin}_t \sum_j w_j |y'_j - t|$, where c is a consistency factor to warrant consistent estimation of σ in case of Gaussian observations, i.e., $c = \operatorname{argmin}_t \mathbb{E} \sum_j w_j |\tilde{y}_j - t|$ for $\tilde{y}_j \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$. c can be obtained empirically for a sufficiently large sample size M , e.g., $M = 10000$, setting $c = \frac{1}{M} \sum_{k=1}^M c_k$, $c_k = \operatorname{argmin}_t \sum_j w_j |y''_{j,k} - t|$, $y''_{j,k} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$.

As to the (finite sample) breakdown point *FSBP* of the weighted median (and at the same time for the scaled weighted MAD), we define $w_j^0 = w_{i,j} / \sum_{j'} w_{j'}$, and for each i define the ordered weights $w_{(j)}^0$ such that $w_{(1)}^0 \geq w_{(2)}^0 \geq \dots \geq w_{(k)}^0$. Then the FSBP in both cases is $k^{-1} \min\{j_0 = 1, \dots, k \mid \sum_{j=1}^{j_0} w_{(j)}^0 \geq k/2\}$ which (for equivariant estimators) can be shown to be the largest possible value. So using weighted medians and MADs, we achieve a decent degree of robustness against outliers. E.g., assume we have 10 observations with weights $5 \times 0.05; 3 \times 0.1; 0.2; 0.25$. Then we need at least three outliers (placed at weights 0.1, 0.2, 0.25, respectively) to produce a breakdown.

5.4 Robustification of the E-step

As indicated, in this step we cannot recur to asymptotics, but rather have to appeal to a theory particularly suited for this recursive setting. In particular, the SO-neighborhoods introduced in (4.3) turn out to be helpful here.

5.4.1 Crucial Optimality Thm

Consider the following optimization problem of reconstructing the ideal observation Y^{id} by means of the realistic/possibly contaminated Y^{re} on an SO-neighborhood.

Minimax-SO problem Minimize the maximal MSE on an SO-neighborhood, i.e., find a Y^{re} -measurable reconstruction f_0 for Y^{id} s.t.

$$\max_{\mathcal{U}} \mathbb{E}_{\text{re}} |Y^{\text{id}} - f(Y^{\text{re}})|^2 = \min_f ! \quad (5.1)$$

The solution is given by

Theorem 5.1 (Minimax-SO) *In this situation, there is a saddle-point $(f_0, P_0^{Y^{\text{di}}})$ for Problem (5.1)*

$$f_0(y) := \mathbb{E} Y^{\text{id}} + H_\rho(D(y)), \quad H_b(z) = z \min\{1, b/|z|\} \quad (5.2)$$

$$P_0^{Y^{\text{di}}}(dy) := \frac{1-r}{r} (|D(y)|/\rho - 1)_+ P^{Y^{\text{id}}}(dy) \quad (5.3)$$

where $\rho > 0$ ensures that $\int P_0^{Y^{\text{di}}}(dy) = 1$ and

$$D(y) = y - \mathbb{E} Y^{\text{id}} \quad (5.4)$$

The value of the minimax risk of Problem (5.1) is

$$\text{tr Cov}(Y^{\text{id}}) - (1-r) \mathbb{E}_{\text{id}} [\min\{|D(Y^{\text{id}})|, \rho\}^2] \quad (5.5)$$

PROOF : See Appendix 7. ////

The optimal procedure in equation (5.2) has an appealing interpretation: It is a compromise between the (unobservable) situations that (a) one observes the ideal Y^{id} , in which case one would use it unchanged and (b) one observes Y^{di} , i.e.; something completely unrelated to Y^{id} , hence one would use the best prediction for Y^{id} (in MSE-sense) without any information, hence the unconditional expectation $\mathbb{E} Y^{\text{id}}$. The decision on how much to tend to case (a) and how much to case (b) is taken according to the (length of the) discrepancy $D(Y^{\text{re}})$ between observed signal Y^{re} and $\mathbb{E} Y^{\text{id}}$. If this length is smaller than ρ , we keep Y^{re} unchanged, otherwise we modify $\mathbb{E} Y^{\text{id}}$ by adding a clipped version of $D(Y^{\text{re}})$.

5.4.2 Robustification of Steps (RN), (E)

In the Girsanov-/change of measure step we recall that the corresponding likelihood ratio here is just

$$\lambda_s := \frac{\sigma^{-1}(\mathbf{x}_{s-1}) \varphi\left(\sigma^{-1}(\mathbf{x}_{s-1})(y_s - f(\mathbf{x}_{s-1}))\right)}{\varphi(y_s)} \quad (5.6)$$

Apparently λ_s can both take values close to 0, and, more dangerously, in particular for small values of $\sigma(\mathbf{x}_{s-1})$, very large values. So bounding the λ_s is crucial to avoid effects like in Figure 3.

A first (non-robust) remedy uses a data driven reference measure, i.e., instead of $\mathcal{N}(0, 1)$, we use $\mathcal{N}(0, \bar{\sigma}^2)$ where $\bar{\sigma}$ is a global scale measure taken over all observations, ignoring time dependence and state-varying σ 's. A robust proposal would take $\bar{\sigma}$ to be the MAD of all observations (tuned for consistency at the normal distribution). This leads to

$$\tilde{\lambda}_s := \frac{\sigma^{-1}(\mathbf{x}_{s-1}) \varphi\left(\sigma^{-1}(\mathbf{x}_{s-1})(y_s - f(\mathbf{x}_{s-1}))\right)}{\bar{\sigma}^{-1} \varphi(\bar{\sigma}^{-1} y_s)} \quad (5.7)$$

Eventually, in both estimation and filtering/prediction, $\bar{\sigma}$ cancels out as a common factor in nominator and denominator, so is irrelevant in the subsequent steps; its mere purpose is to stabilize the terms in numeric aspects.

To take into account time dynamics in our robustification, we want to use Theorem 5.1, but to this end, we need second moments, which for λ_s need not exist. So instead, we apply the theorem to $Y^{\text{id}} = \sqrt{\tilde{\lambda}_s}$, which means that $\tilde{\lambda}_s = (Y^{\text{id}})^2$ is robustified by

$$\bar{\lambda}_s = (\mathbb{E}_{\text{id}} \sqrt{\tilde{\lambda}_s} + H_b(\sqrt{\tilde{\lambda}_s} - \mathbb{E}_{\text{id}} \sqrt{\tilde{\lambda}_s}))^2 \quad (5.8)$$

for $H_b(x) = x \min\{1, b/|x|\}$. Clipping height b in turn is chosen such that $\mathbb{E} \bar{\lambda}_s = \alpha$, $\alpha = 0.95$ for instance. As in the ideal situation $\mathbb{E} \lambda_s = 1$, in a last step with a consistency factor c'_s determined similarly to c_i in the initialization step for the weighted MADs, we pass over to $\bar{\lambda}_s^0 = c_s \bar{\lambda}_s$ such that $\mathbb{E} \bar{\lambda}_s^0 = 1$.

Similarly, in the remaining parts of the E-step, for each of the filtered processes generically denoted by G and the filtered one by \hat{G} , we could replace \hat{G} by

$$\tilde{G} = \mathbb{E}_{\text{id}} \hat{G} + H_b(\hat{G} - \mathbb{E}_{\text{id}} \hat{G}) \quad (5.9)$$

for G any of J_k^{ji} , O_k^i , and $T_k^i(f)$ and again suitably chosen b .

It turns out though, that it is preferable to pursue another route. The aggregates $T_k^i(f)$ are used in the M-step in (3.10), but for a robustification of this step, it is crucial to be able to attribute individual influence to each of the observations, so instead we split up the terms of the filtered neg-loglikelihood into summands $w_{i,j}/\hat{O}_k^i[(y_j - f_i)^2/\sigma_i^2 + \log \sigma_i]$ for $j = 1, \dots, k$, and hence we may skip a robustification of $T_k^i(f)$. Similarly, as J_t^{ji} , O_t^i are filtered observations of multinomial-like variables, a robustification is of limited use, as any contribution of a single observation to these variables can at most be of absolute value 1, so is bounded anyway. Hence in the E-step, we only robustify λ_s . The splitting up the aggregates $T_k^i(f)$ into summands amounts to giving up strict recursivity, as for k observations in one batch, one now has to store the values $w_{i,j}/\hat{O}_k^i$ for $j = 1, \dots, k$, and building up from $j = 1$, at observation time $j = j_0$ within the batch, we construct $w_{i,j;j_0}$, $j = 1, \dots, j_0$ from the values $w_{i,j;j_0-1}$, $j = 1, \dots, j_0 - 1$, so we have a growing triangle of weight values. This would lead to increasing memory requirements, if we had not chosen to work in batches of fixed length k , which puts an upper bound onto memory needs.

5.5 Robustification of the (M)-Step

As mentioned before, contrast to the (E)-step, in this estimation step, we may work with classical gross error neighborhoods (4.1) and with the standard i.i.d. setting.

5.5.1 Shrinking Neighborhood Approach

By Bienaymé, variance then usually is $O(1/n)$ for sample size n , while for robust estimators, the maximal bias is proportional to the neighborhood radius ε . Hence unless ε is appropriately scaled in n , for growing n , bias will dominate eventually for growing n . This is avoided in the shrinking neighborhood approach by setting $\varepsilon = \varepsilon_n = r/\sqrt{n}$ for some $r \geq 0$, compare Rieder [1994]. Kohl et al. [2010] sets $\varepsilon = \varepsilon_n = r/\sqrt{n}$ for some initial radius $r \in [0, \infty)$. One could see this shrinking as indicating that with growing n , diligence is increasing so the rate of outliers is decreasing. This is perhaps overly optimistic. Another interpretation is that the severeness of the robustness problem with 10% outliers at sample size 100 should not be compared with the one with 10% outliers at sample size 10000 but rather with the one with 1% outliers at this sample size.

In this shrinking neighborhood setting, with mathematical rigor, optimization of the robust criteria can be deferred to the respective IFs, i.e. instead of determining the IF of a given procedure, we construct a procedure to a given (optimally-robust) IF.

This is achieved by the concept of *asymptotically linear estimators* (ALEs), as it arises canonically in most proofs of asymptotic normality: In a smooth (L_2 -differentiable) parametric model $\mathcal{P} = \{P_\theta, \theta \in \Theta\}$ for i.i.d. observations $X_i \sim P_\theta$ with open parameter domain $\Theta \subset \mathbb{R}^d$ based on the scores³ Λ_θ and its finite Fisher information $\mathcal{I}_\theta = E_\theta \Lambda_\theta \Lambda_\theta^T$, we define the set $\Psi_2(\theta)$ of influence functions as the subset of $L_2^d(P_\theta)$ consisting of square integrable functions ψ_θ with d coordinates with $E_\theta \psi_\theta = 0$ and $E_\theta \psi_\theta \Lambda_\theta^T = \mathbb{I}_d$ where \mathbb{I}_d is the d -dimensional unit matrix. Then a sequence of estimators $S_n = S_n(X_1, \dots, X_n)$ is called an ALE if

$$S_n = \theta + \frac{1}{n} \sum_{i=1}^n \psi_\theta(X_i) + o_{P_\theta^n}(n^{-1/2}) \quad (5.10)$$

for some influence function $\psi_\theta \in \Psi_2(\theta)$. In the sequel we fix the true $\theta \in \Theta$ and suppress it from notation where clear from context.

In particular, the MLE usually has influence function $\psi^{\text{MLE}} = \mathcal{I}^{-1} \Lambda$, while most other common estimators also have a representation (5.10) with a different ψ .

For given IF ψ we may construct an ALE $\hat{\theta}_n$ with ψ as IF by a one-step construction, often called one-step-reweighting: To given starting estimator θ_n^0 such that $R_n^0 = \theta_n^0 - \theta = o_{P_\theta^n}(n^{-1/4+0})$ we define

$$\hat{\theta}_n = \theta_n^0 + \frac{1}{n} \sum_{j=1}^n \psi_{\theta_n^0}(X_j) \quad (5.11)$$

Then indeed $\hat{\theta}_n = \theta + \frac{1}{n} \sum_{j=1}^n \psi_\theta(X_j) + R_n$ and $R_n = o_{P_\theta^n}(n^{-1/2})$, i.e., $\hat{\theta}_n$ forgets about θ_n^0 as to its asymptotic variance and GES; however its breakdown point is inherited from θ_n^0 once Θ is unbounded and ψ is bounded. Hence for the starting estimator, we seek for θ_n^0 with high breakdown point.

³Usually Λ_θ is the logarithmic derivative of the density w.r.t. the parameter, i.e., $\Lambda_\theta(x) = \partial/\partial\theta \log p_\theta(x)$.




Figure 4: Influence function of the MBRE at $\mathcal{N}(0,1)$; left panel: location part; right panel: scale part.

For a more detailed account on this approach, see Rieder [1994].

5.5.2 Shrinking Neighborhood Approach With Weighted Observations

As mentioned, coming from the E-step, not all observations y_j are equally likely to contribute to state i , hence we are in a situation with weighted observations, where we may pass over to normed weights $w_{i,j}^0 = w_{i,j} / \sum_{j'} w_{i,j'}$ summing up to 1.

Suppressing state index i from notation here, with these weights and θ_n^0 the vector of weighted median and scaled weighted MAD for state i , (5.11) becomes

$$\hat{\theta}_n = \theta_n^0 + \sum_{j=1}^n w_j^0 \psi_{\theta_n^0}(y_j) \quad (5.12)$$

for $\hat{\theta}_n$ again a two-dimensional ALE with location and scale coordinate and $\psi_{\theta}(y) = \sigma \psi((y - f)/\sigma)$, at $\theta = (f, \sigma)^\tau$, is the IF of the MBRE in the one-dimensional Gaussian location and scale model at $\mathcal{N}(0,1)$; i.e.,

$$\psi(y) = bY(y)/|Y(y)|, \quad Y(y) = (y, A(y^2 - 1) - a), \quad (5.13)$$

with numerical values for A, a, b up to four digits taken from R package **RobLox**, Kohl [2012], being

$$A = 0.7917, \quad a = -0.4970, \quad b = 1.8546, \quad (5.14)$$

Influence function ψ is illustrated in Figure 4.

To warrant positivity of σ and to maintain a high breakdown point even in the presence of inliers (driving σ to essentially 0), for the scale component, instead of (5.12) we use the asymptotically equivalent form

$$\hat{\sigma}_n = \sigma_n^0 \exp \left[\sum_{j=1}^n w_j^0 \psi_{\text{scale}}((y_j - \mu_n^0)/\sigma_n^0) \right]. \quad (5.15)$$

Using the MBRE- ψ from (5.13) and (5.14) on first glance could be seen as overly cautious. Detailed simulation studies, compare e.g. Kohl and Deigner [2010], show that for our typical batch lengths of 10–20, the MBRE also is near to optimal in the sense of Rieder et al. [2008] in the situation where nothing is known about the true outlier rate (including, of course the situation where no outliers at all occur).

5.5.3 Robustification of Steps (M1) & (M2)

Now, we derive robust estimators of the model parameters f_i and σ_i , i.e., we justify passage to weighted ALEs as in (5.12). In particular we specify the weights $w_j^0 = w_{i,j}^0$ therein.

Recall, that the M1-step in the classical algorithm gives the optimal parameter estimates stated in Theorem 3.1. We now build ALEs, which can be achieved, when the MLEs of the parameters f_i and σ_i are stated as weighted sums of the observations y_k .

Theorem 5.2 *With*

$$w_{i,l}^0 = \frac{\langle \hat{\mathbf{x}}_{l-1}, \mathbf{e}_i \rangle}{\eta(O^{(i)})_k} \quad (5.16)$$

the optimal parameter estimates \hat{f}_i and $\hat{\sigma}_i$ are given by

$$\hat{f}_i = \sum_{l=1}^k w_{i,l}^0 y_l \quad (5.17)$$

$$\hat{\sigma}_i^2 = \sum_{l=1}^k w_{i,l}^0 (y_l - \hat{f}_i)^2. \quad (5.18)$$

PROOF : To find the optimal estimate for f consider

$$\Lambda_k^* := \prod_{l=1}^k \lambda_l^* \text{ with } \lambda_l^* := \exp\left(\frac{(y_l - \langle f, \mathbf{x}_{l-1} \rangle)^2 - (y_l - \langle \hat{f}, \mathbf{x}_{l-1} \rangle)^2}{2\langle \sigma, \mathbf{x}_{l-1} \rangle^2}\right).$$

Up to constants irrelevant for optimization, the filtered log-likelihood is then

$$\mathbb{E}[\ln(\Lambda_k^*) \mid \mathcal{F}_k^y] = \sum_{i=1}^N \sum_{l=1}^k w_{i,l}^0 (y_l - \hat{f}_i)^2 \quad (5.19)$$

Maximising the log-likelihood $\mathbb{E}[\ln(\Lambda_k^*) \mid \mathcal{F}_k^y]$ in \hat{f}_i hence leads to the optimal parameter estimate

$$\sum_{l=1}^k \langle \mathbf{x}_{l-1}, \mathbf{e}_i \rangle (2y_l \hat{f}_i - \hat{f}_i^2) = 0 \quad \implies \quad \hat{f}_i(k) = \sum_{l=1}^k w_{i,l}^0 y_l$$

In an analogue way, for σ_i we define

$$\Lambda_k^+ := \prod_{l=1}^k \lambda_l^+ \text{ with } \lambda_l^+ := \frac{\langle \sigma, \mathbf{x}_{l-1} \rangle}{\langle \hat{\sigma}, \mathbf{x}_{l-1} \rangle} \exp\left(\frac{(y_l - \langle f, \mathbf{x}_{l-1} \rangle)^2}{2\langle \sigma, \mathbf{x}_{l-1} \rangle^2} - \frac{(y_l - \langle f, \mathbf{x}_{l-1} \rangle)^2}{2\langle \hat{\sigma}, \mathbf{x}_{l-1} \rangle^2}\right).$$

and hence, again up to irrelevant terms

$$\mathbb{E}[\ln(\Lambda_k^+) \mid \mathcal{F}_k^y] = \sum_{i=1}^N \sum_{l=1}^k \langle \hat{\mathbf{x}}_{l-1}, \mathbf{e}_i \rangle \left[\frac{(y_l - \hat{f}_i)^2}{2\hat{\sigma}_i^2} + \log(\sigma_i) \right] \quad (5.20)$$

From this term, which has to be minimised for $\hat{\sigma}_i$ we get

$$\sum_{l=1}^k \langle \mathbf{x}_{l-1}, \mathbf{e}_i \rangle \left[-\frac{1}{2\hat{\sigma}_i^3} (y_l - \hat{f}_i)^2 + \frac{1}{\hat{\sigma}_i} \right] = 0 \quad \implies \quad \hat{\sigma}_i^2 = \sum_{l=1}^k w_{i,l}^0 (y_l - \hat{f}_i)^2$$

Note that (5.20) takes its minimum at the same place as

$$\ln(\Lambda_k^{++}) = \sum_{i=1}^N \sum_{l=1}^k w_{i,l}^0 [(y_l - \hat{f}_i)^2 - \sigma_i^2]^2 \quad (5.21)$$

////

For the robustification of the parameter estimation (step M1) we now distinguish two approaches. The first robustification is utilized in the first run over the first batch of data and is therefore called the *initialization* step M1. The robust estimates of the parameters from the second batch onwards are then achieved through a weighted ALE.

Theorem 5.3 *The robust parameter estimates for the model parameters f_i and σ_i in the (1) initialization and (2) all following batches are given by*

1. *Replacing, for initialization, the squares by absolute values in (5.19) and in (5.21), \hat{f}_i and $\hat{\sigma}_i$ are the weighted median and scaled weighted MAD, respectively, of the y_l , $l = 1, \dots, k$ with weights $(w_{i,l}^0)_l$.*
2. *For further batches, the weighted MBRE is obtained as a one-step construction with the parameter estimate (f_i^0, σ_i^0) from the previous batch as starting estimator and with IF $\psi = (\psi_{\text{loc}}, \psi_{\text{scale}})$ from (5.13), (5.14), i.e.,*

$$\hat{f}_i = f_i^0 + \sigma_i^0 \sum_{l=1}^k w_{i,l}^0 \psi_{\text{loc}}((y_l - f_i^0)/\sigma_i^0) \quad (5.22)$$

$$\hat{\sigma}_i = \sigma_i^0 \exp\left(\sum_{l=1}^k w_{i,l}^0 \psi_{\text{scale}}((y_l - f_i^0)/\sigma_i^0)\right) \quad (5.23)$$

PROOF :

1. Initialization: With absolute values instead of squares, (5.19) becomes

$$\hat{f}_i = \underset{f_i}{\operatorname{argmin}} \sum_{l=1}^k w_{i,l}^0 |y_l - f_i|$$

Now if $w_{i,l}^0$ is constant in l , this leads to the empirical median as unique minimizer justifying the name. For the scaled weighted MAD, the argument parallels the previous one, leading to consistency factor $c_i = \Phi(3/4)$ for Φ the cdf of $\mathcal{N}(0, 1)$.

2. M1 in further batches: Apparently, by definition, $(\hat{f}_i, \hat{\sigma}_i)$ is an ALE, once we show that ψ is square integrable, $E(\psi) = 0$, $E(\psi\Lambda') = \mathbb{I}_2$. The latter two properties can be checked numerically, while by boundedness square integrability is obvious. In addition it has the necessary form of an MBRE in the i.i.d. setting as given in [Rieder, 1994, Thm 5.5.1]. To show that this also gives the MBRE in the context of weighted observations, we would need to develop the theory of ALEs for triangular schemes similar to the one in the Lindeberg Feller Theorem. This has been done, to some extent in [Ruckdeschel, 2001, Section 9]. In particular, for each state i , we have to assume a Noether condition excluding observations overly influential for parameter estimation in this particular state, i.e.,

$$\lim_{k \rightarrow \infty} \max_{l=1, \dots, k} (w_{i,l;k}^0)^2 / \sum_{j=1}^k (w_{i,j;k}^0)^2 = 0 \quad (5.24)$$

We do not work this out in detail here, though.

////




Figure 5: Robust parameter estimates for monthly MSCI returns between 1994 and 2009—analogue to Figure 1

Consider again our filtering algorithm and recall, that the filter runs over the data set in batches of roughly ten to fifty data point. To determine the ALE for our parameters, we have to calculate the weights $w_{i,l}^0 = \langle \hat{\mathbf{x}}_{l-1}, \mathbf{e}_i \rangle / \hat{O}_k^i$. Therefore, our algorithm has to know all values of $\hat{\mathbf{x}}_l$ from 1 to k in each batch. With this, our robustification of the algorithm cannot obtain the same recursiveness as the classical algorithm. However, since we only have to determine and save the estimates of \mathbf{x}_l in each batch, the algorithm still is numerically efficient, the additional costs are low. In general, the ALEs are fastly computed robust estimators, which lead in our case to a fast and, over batches, recursive algorithm.

The additional computational burden to store all the weights $w_{i,l}^0$ arising in the robustification of the M1-step is more than paid off by the additional benefits they offer for diagnostic purposes beyond the mere EM-algorithm: They tell us which of the observations, due to their likelihood to be in state i carry more information on the respective parameters f_i and σ_i than others. The same goes for the terms $\psi_\theta(y_l)$ which capture the individual information of observation y_l for the respective parameters. Even more though, the coordinates of $\psi_\theta(y_j)/|\psi_\theta(y_l)|$ tell us how much of the information in observation y_l is used for estimating f_i and how much for σ_i . In addition the function $y \mapsto w_{i,l}^0 \psi_\theta(y)$ can be used for sensitivity analysis, telling us what happens to the parameter estimates for small changes in observation y . Finally, using the unclipped, classically optimal IF of the MLE, but evaluated at the robustly estimated parameters, we may identify outliers not fitting to the “usual” states.

6 Implementation and Simulation

The classical algorithm as well as the robust version are implemented in R; we plan to release the code in form of a contributed package on CRAN at a later stage. The implementation builds up on, respectively uses contributed packages `RobLox` and `mclust`. At the time of writing we are preparing a thorough simulation study to explore our procedure in detail and in a quantitative way. For the moment, we restrict ourselves to assess the procedure in a qualitative way, illustrating how it can cope with a situation like in Figure 3.

In Figure 5, we see the paths of the robust parameter estimates for f , σ , and Π ; due to the new initialization procedure, the estimates—in particular those for Π —differ a little from those of Figure 1. Still, all the estimators behave very reasonable and are not too far from the classical ones.

In the outlier situations from Figures 2 and 3, illustrated in Figures 6, the estimates for f and σ remain stable at large as desired. The estimates for Π however do get irritated, essentially flagging out one state as outlier state. Some more work remains to be done to better understand this and to see how to avoid this.

Aside from this, our algorithm already achieves its goals; in particular, our procedure never breaks down—contrary to the classical one.

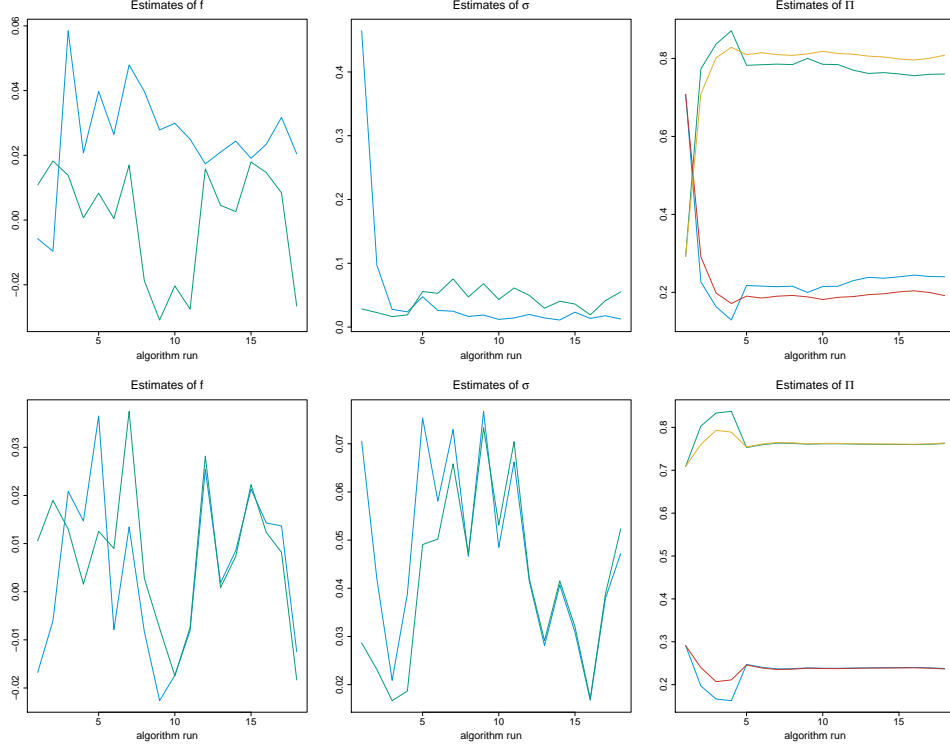


Figure 6: Robust parameter estimates for monthly MSCI with planted outliers—analogue to Figures 2 and 3

7 Conclusion

In financial applications, we often have to consider the case of outliers in our data set, which can occur from time to time e.g., due to either wrong values in the financial database or unusual peaks or lows in volatile markets. Conventional parameter estimation methods cannot handle these specific data characteristics well.

Contribution of this paper: Our contribution to this issue is two fold:

First, we analyse step by step the general filter-based EM-algorithm for HMMs by Elliott [1994] and highlight, which problems can occur in case of extreme values. We extend the classical algorithm by a new technique to find initial values, taking into account the N —state setting of the HMM. In addition, for numerical reasons we use a data-driven reference measure instead of the standard normal distribution.

Second, we have proposed a full robustification of the classical EM-algorithm. Our robustified algorithm is stable w.r.t. outliers in the observation process and is still able to estimate processes of the Markov chain as well as optimal parameter estimates with acceptable accuracy. The robustification builds up on concepts from robust statistics like SO-optimal filtering and asymptotic linear estimators. Due to the non-iid nature of the observations as apparent from the non-uniform weights $w_{i,l}^0$ attributed to the observations, these concepts had to be generalized for this situation, leading to weighted medians, weighted MADs, weighted ALEs. Similarly, the SO-optimal filtering (with focus on state reconstruction) is not directly applicable for robustifying the Radon-Nikodym terms λ_s , where we (a) had to clean the “observations” themselves and (b) had to pass over to $\sqrt{\lambda_s}$ for integrability reasons.

Our robust algorithm is computationally efficient. Although complete recursivity cannot be obtained, the algorithm runs over batches and keeps its recursivity there additionally storing the filtered values of the Markov chain. This additional burden is outweighed though by the benefits of these weights and influence function terms for diagnostic purposes. As in the original algorithm, the model parameters, which are guided by the state of the Markov chain, are updated after each batch, using a robust ALE however. The robustification therefore keeps the characteristic of the algorithm, that new information, which arises in the observation process, is included in the recent parameter update—there is no forward-backward loop.

The forecasts of asset prices, which are obtained through the robustified parameter estimates, can be utilized to make investment decisions in asset allocation problems.

To sum up, our forecasts are robust against additive outliers in the observation process and able to handle switching regimes occurring in financial markets.

Outlook: It is pretty obvious how to generalize our robustification to a multivariate setting: The E-step is not affected by multivariate observations, and the initialization technique using Gaussian Mixture Models ideas already is available in multivariate settings. Respective robust multivariate scale and location estimators for weighted situations still have to be implemented, though, a candidate being a weighted variant of the (fast) MCD-estimator, compare Rousseeuw and Leroy [1987], Rousseeuw and van Driessen [1999].

Future work will hence translate our robustification to a multivariate setting to directly apply the algorithm to asset allocation problems for portfolio optimisation. Furthermore, investment strategies shall be examined within this robust HMM setting to enable investors a view on their portfolio, which includes possible outliers or extreme events. The implementation of the algorithms shall be part of an R package, including a thorough simulation study of the robustified algorithm and its application in portfolio optimisation.

Finally, an automatic selection criterion for the number of states to retain would be desirable which is a question of model selection, where criteria like BIC have still to be adopted for robustness.

Acknowledgement

Financial support for C. Erlwein from Deutsche Forschungsgemeinschaft (DFG) within the project “Regimeswitching in zeitstetigen Finanzmarktmodellen: Statistik und problemspezifische Modellwahl” (RU-893/4-1) is gratefully acknowledged.

A Proofs

Proof to Theorem 5.1

(1) Let us solve $\max_{\partial\mathcal{U}} \min_f [\dots]$ first, which amounts to $\min_{\partial\mathcal{U}} \mathbb{E}_{\text{re}}[|\mathbb{E}_{\text{re}}[Y^{\text{id}}|Y^{\text{re}}]|^2]$. For fixed element $P^{Y^{\text{di}}}$ assume a dominating σ -finite measure μ , i.e., $\mu \gg P^{Y^{\text{di}}}$, $\mu \gg P^{Y^{\text{id}}}$; this gives us a μ -density $q(y)$ of $P^{Y^{\text{di}}}$. Determining the joint (real) law $P^{Y^{\text{id}}, Y^{\text{re}}}(d\tilde{y}, dy)$ as

$$P(Y^{\text{id}} \in A, Y^{\text{re}} \in B) = \int \mathbb{I}_A(\tilde{y}) \mathbb{I}_B(y) [(1-r) \mathbb{I}(\tilde{y} = y) + r q(y)] p^{Y^{\text{id}}}(\tilde{y}) \mu(d\tilde{y}) \mu(dy) \quad (\text{A.1})$$

we deduce that $\mu(dy)$ -a.e.

$$\mathbb{E}_{\text{re}}[Y^{\text{id}}|Y^{\text{re}}=y] = \frac{r q(y) \mathbb{E} Y^{\text{id}} + (1-r) y p^{Y^{\text{id}}}(y)}{r q(y) + (1-r) p^{Y^{\text{id}}}(y)} =: \frac{a_1 q(y) + a_2(y)}{a_3 q(y) + a_4(y)} \quad (\text{A.2})$$

Hence we have to minimize

$$F(q) := \int \frac{|a_1 q(y) + a_2(y)|^2}{a_3 q(y) + a_4(y)} \mu(dy)$$

in $M_0 = \{q \in L_1(\mu) \mid q \geq 0, \int q d\mu = 1\}$. To this end, we note that F is convex on the non-void, convex cone $M = \{q \in L_1(\mu) \mid q \geq 0\}$ so, for some $\tilde{\rho} \geq 0$, we may consider the Lagrangian

$$L_{\tilde{\rho}}(q) := F(q) + \tilde{\rho} \int q d\mu$$

for some positive Lagrange multiplier $\tilde{\rho}$. Pointwise minimization in y of $L_{\tilde{\rho}}(q)$ gives

$$q_s(y) = \frac{1-r}{r} (|D(y)| / s - 1)_+ p^Y(y)$$

for some constant $s = s(\tilde{\rho}) = (|\mathbb{E} Y^{\text{id}}|^2 + \tilde{\rho}/r)^{1/2}$, Pointwise in y , \hat{q}_s is antitone and continuous in $s \geq 0$ and $\lim_{s \rightarrow 0[\infty]} q_s(y) = \infty[0]$, hence by monotone convergence,

$$H(s) = \int \hat{q}_s(y) \mu(dy)$$

too, is antitone and continuous and $\lim_{s \rightarrow 0[\infty]} H(s) = \infty[0]$. So by continuity, there is some $\rho \in (0, \infty)$ with $H(\rho) = 1$. On M_0 , $\int q d\mu = 1$, but $\hat{q}_\rho = q_{s=\rho} \in M_0$ and is optimal on $M \supset M_0$ hence it also minimizes F on M_0 . In particular, we get representation (5.3) and note that, independently from the choice of μ , the least favorable $P_0^{Y^{\text{di}}}$ is dominated according to $P_0^{Y^{\text{di}}} \ll P^{Y^{\text{id}}}$, i.e.; non-dominated $P^{Y^{\text{id}}}$ are even easier to deal with.

As next step we show that

$$\max_{\partial \mathcal{U}} \min_f [\dots] = \min_f \max_{\partial \mathcal{U}} [\dots] \quad (\text{A.3})$$

To this end we first verify (5.2) determining $f_0(y)$ as $f_0(y) = E_{\text{re}; \hat{P}}[X|Y^{\text{re}} = y]$. Writing a sub/superscript “re; \hat{P} ” for evaluation under the situation generated by $P = P^{Y^{\text{di}}}$ and \hat{P} for $P_0^{Y^{\text{di}}}$, we obtain the the risk for general P as

$$\begin{aligned} \text{MSE}_{\text{re}; P}[f_0(Y^{\text{re}}, P)] &= (1-r) E_{\text{id}} |Y^{\text{id}} - f_0(Y^{\text{id}})|^2 + r \text{tr Cov } Y^{\text{id}} + \\ &\quad + r E_P \min(|D(Y^{\text{di}; q})|^2, \rho^2) \end{aligned} \quad (\text{A.4})$$

This is maximal for any P that is concentrated on the set $\{|D(Y^{\text{di}; q})| > \rho\}$, which is true for \hat{P} . Hence (A.3) follows, as for any contaminating P

$$\text{MSE}_{\text{re}; P}[f_0(Y^{\text{re}}; P)] \leq \text{MSE}_{\text{re}; \hat{P}}[f_0(Y^{\text{re}}; \hat{P})]$$

Finally, we pass over from $\partial \mathcal{U}$ to \mathcal{U} : Let f_r , \hat{P}_r denote the components of the saddle-point for $\partial \mathcal{U}(r)$, as well as $\rho(r)$ the corresponding Lagrange multiplier and w_r the corresponding weight, i.e., $w_r = w_r(y) = \min(1, \rho(r) / |D(y)|)$. Let $R(f, P, r)$ be the MSE of procedure f at the SO model $\partial \mathcal{U}(r)$ with contaminating $P^{Y^{\text{di}}} = P$. As can be seen from (5.3), $\rho(r)$ is antitone in r ; in particular, as \hat{P}_r is concentrated on $\{|D(Y)| \geq \rho(r)\}$ which for $r \leq s$ is a subset of $\{|D(Y)| \geq \rho(s)\}$, we obtain

$$R(f_s, \hat{P}_s, s) = R(f_s, \hat{P}_r, s) \quad \text{for } r \leq s$$

Note that $R(f_s, P, 0) = R(f_s, Q, 0)$ for all P, Q —hence passage to $\tilde{R}(f_s, P, r) = R(f_s, P, r) - R(f_s, P, 0)$ is helpful—and that

$$\text{tr Cov } Y^{\text{id}} = E_{\text{id}} \left[\text{tr Cov}_{\text{id}}[Y^{\text{id}}|Y^{\text{id}}] + |D(Y^{\text{id}})|^2 \right] \quad (\text{A.5})$$

Abbreviate $\bar{w}_s(Y^{\text{id}}) = 1 - (1 - w_s(Y^{\text{id}}))^2 \geq 0$ to see that

$$\begin{aligned} \tilde{R}(f_s, P, r) &= r \left\{ E_{\text{id}} \left[|D(Y^{\text{id}})|^2 \bar{w}_s(Y^{\text{id}}) \right] + E_P \min(|D(Y^{\text{id}})|, \rho(s))^2 \right\} \leq \\ &\leq r \left\{ E_{\text{id}} \left[|D(Y^{\text{id}})|^2 \bar{w}_s(Y^{\text{id}}) \right] + \rho(s)^2 \right\} = \tilde{R}(f_s, \hat{P}_r, r) < \tilde{R}(f_s, \hat{P}_s, s) \end{aligned}$$

Hence the saddle-point extends to $\mathcal{U}(r)$; in particular the maximal risk is never attained in the interior $\mathcal{U}(r) \setminus \partial \mathcal{U}(r)$. (5.5) follows by plugging in the results.

////

References

- ANG A, BEKAERT G. (2002): International asset allocation with regime shifts. *Review of Financial Studies* **15**:1137 - 1187.
- CAI, J. (1994): A Markov model of switching-regime ARCH. *Journal of Business & Economic Statistics* **12**: 309–316.
- ELLIOTT R.J., AGGOUN L., AND MOORE J.B. (1995): *Hidden Markov models: Estimation and Control*. Applications of Mathematics vol. 29. Springer: New York.
- Elliott R.J. (1994): Exact adaptive filters for Markov chains observed in Gaussian noise. *Automatica* **30**: 1399–1408.
- ELLIOTT R.J. AND VAN DER HOEK J. (1997): An application of hidden Markov models to asset allocation problems. *Finance and Stochastics* **1**: 229–238.
- ELLIOTT R.J, HINZ J. (2003): A method for portfolio choice. *Applied Stochastic Models in Business and Industry* **19**: 1–11.
- ERLWEIN C., MAMON R., AND DAVISON M. (2009): An examination of HMM-based investment strategies for asset allocation. *Applied stochastic models in business and industry* **27(3)**: 204–221.

- FERNHOLZ, L.T. (1983): Von Mises Calculus for Statistical Functionals. *Lecture Notes in Statistics*, vol. 19, Springer.
- FOX, A.J. (1972): Outliers in time series. *J. R. Stat. Soc., Ser. B*, **34**: 350–363.
- FRALEY, C. AND RAFTERY, A.E. (2002): Model-based Clustering, Discriminant Analysis and Density Estimation. *Journal of the American Statistical Association* **97**: 611–631.
- FRALEY, C., RAFTERY, A.E., MURPHY, T.B., AND SCRULLA, L. (2012): *mclust* Version 4 for R: Normal Mixture Modeling for Model-Based Clustering, Classification, and Density Estimation. Technical Report No. 597, Department of Statistics, University of Washington.
- GRAY, S.F. (1996): Modeling the conditional distribution of interest rates as a regime-switching process. *Journal of Financial Economics* **42**: 27–62.
- GUIDOLIN M, TIMMERMAN A. (2007): Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control* **31**: 3503–3544.
- HAMILTON, J.D. (1989): A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica: Journal of the Econometric Society* **57**: 357–384.
- HAMPEL, F.R. (1968): *Contributions to the theory of robust estimation*. Dissertation, University of California, Berkeley, CA.
- HAMPEL, F. R., RONCHETTI, E. M., ROUSSEEUW, P. J. AND STAHEL, W. A. (1986): *Robust statistics. The approach based on influence functions*. Wiley.
- HUBER, P. J. (1981): *Robust statistics*, Wiley.
- KOHL, M. (2012). RobLox: Optimally robust influence curves and estimators for location and scale. R package version 0.8.2. URL <http://robast.r-forge.r-project.org/>.
- KOHL, M. AND DEIGNER, H.P. Preprocessing of gene expression data by optimally robust estimator. *BMC Bioinformatics*, **11**: 583.
- KOHL, M., RIEDER, H., AND RUCKDESCHEL, P. (2010): Infinitesimally Robust Estimation in General Smoothly Parametrized Models. *Stat. Methods Appl.*, **19**: 333–354.
- KNIGHT, F.H. (1921): *Risk, Uncertainty, and Profit*. Boston: Houghton Mifflin.
- MAMON, R., ERLWEIN, C. AND GOPALUNI, B. (2008): Adaptive signal processing of asset price dynamics with predictability analysis. *Information Sciences*, **178**: 203–219.
- MARKOWITZ, H. (1952): Portfolio Selection. *The Journal of Finance*, **7**(1): 77–91.
- MARONNA, R. A., MARTIN, R. D. AND YOHAI, V. J. (2006): *Robust Statistics: Theory and Methods*. Wiley.
- RIEDER, H. (1994): *Robust Asymptotic Statistics*. Springer.
- RIEDER, H., KOHL, M. AND RUCKDESCHEL, P. (2008): The cost of not knowing the radius. *Statistical Methods & Applications* 17(1), 13–40.
- ROUSSEEUW, P.J. AND LEROY, A.M. (1987): *Robust Regression and Outlier Detection*. Wiley.
- ROUSSEEUW, P.J. AND VAN DRIESSEN, K. (1999): A fast algorithm for the minimum covariance determinant estimator. *Technometrics* **41**: 212–223.
- RUCKDESCHEL, P. (2001): *Ansätze zur Robustifizierung des Kalman Filters*. Bayreuther Mathematische Schriften, Vol. 64.
- RUCKDESCHEL, P. AND HORBENKO, N. (2010): Robustness Properties of Estimators in Generalized Pareto Models. Technical Report ITWM N°182, http://www.itwm.fraunhofer.de/fileadmin/ITWM-Media/Zentral/Pdf/Berichte_ITWM/2010/bericht_182.pdf.

- SASS J. AND HAUSSMANN U.G. (2004): Optimizing the terminal wealth under partial information: The drift process as a continuous time Markov chain. *Finance and Stochastics* **8**: 553–577.
- ZAKAI, M. (1969): On the optimal filtering of diffusion processes. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, **11**: 230–243.